

THE DESIGN AND APPLICATION OF BEE2 —A HIGH-END RECONFIGURABLE COMPUTING SYSTEM

Chen Chang, John Wawrzynek, Robert W. Brodersen
{chenzh, johnw, rb}@eecs.berkeley.edu
Department of Electrical Engineering and Computer Sciences
University of California, Berkeley

Abstract

This paper summarizes our effort to design and construct a high-end reconfigurable computer (HERC) system based solely on field programmable gate arrays (FPGAs) as the processing elements. FPGAs offer many important potential advantages over conventional microprocessors and digital signal processors (DSP), such as flexible arithmetic precision, higher computational density per unit silicon area, and lower power consumption. The programmable interconnect structure unique to FPGA technology makes it possible to tailor a HERC system, such as our BEE2 system, on a per-problem basis to best take advantage of task specific data-flow, memory access patterns, and node-to-node communication patterns.

Our BEE2 project is a coordinated attack on the elements needed to demonstrate a practical, cost-effective, high-end reconfigurable computer: the design of a processing module to be used as the building block for a family of high-end reconfigurable computers; the development of several programming models; and the demonstration of the efficiency of the machine on a set of demanding applications, ranging from high-performance digital signal processing and communication systems to traditional scientific computing. On selected DSP applications, BEE2 can provide over 100 times more computing throughput than a microprocessor-based system with similar power consumption and cost.

There are several computationally intensive problems central to the research objectives of BWRC that we are using as an application benchmark set and design drivers for the specification of the BEE2 machine architecture and its associated software mapping tools. These applications fall into four broad categories: high-performance real-time digital signal processing, emulation and design of novel wireless communications systems, real-time scientific computation and simulation, and acceleration of computer aided-design (CAD) tools.

Due to the diverse application domains targeted by the BEE2 system, any single programming model would not be optimal for all applications; hence the need for domain specific programming models that can fully exploit the computing power of the BEE2 system. Currently the most mature programming model for the BEE2 system is the synchronous data flow model for DSP and communication applications. Commercial tools, including Mathworks Matlab/Simulink, Xilinx System Generator, along with automation tools developed at BWRC, provide automatic mapping from high-level block diagrams and state machine specifications to FPGA configurations. This programming model and tool flow has proven very successful on a variety of projects at BWRC, particularly in the areas of DSP and other datapath intensive streaming applications. To extend this model to support BEE2 specific hardware, stream-based design abstractions are currently being developed for external DRAMs and global communication networks.

We have completed the design and fabrication of a compute module comprising five Xilinx XC2VP70 FPGAs, 20 DRAM DIMMs, and 18 off-module 10Gbit/s Infiniband/Ethernet connections, shown below in figure 1. This module has peak performance in the 1-2 TeraOp/s

range (integer operations), and forms the basic building block for larger systems, scalable from 1 to 100's of modules.

To date, our most extensive application development has been in collaboration with the SETI@HOME, SERENDIP project at UC Berkeley Space Science Laboratory and the UC Berkeley Radio Astronomy Laboratory. We have successfully demonstrated an 800MHz billion-channel spectrometer using the BEE2 system on a single antenna. We have analyzed the performance of our FPGA-based approach on this and other Radio Astronomy applications. In terms of computational throughput per chip, the FPGAs in the BEE2 system outperform a 720MHz DSP by a factor of 10 to 34, a 1GHz (90nm) DSP by a factor of 7 to 25, and the latest Pentium-4 by a factor of 4 to 13. In terms of power efficiency, the XC2VP70 FPGA delivers 72% to 106% more throughput on 16bit operations comparing to the DSPs, and more than 11 times on 4-bit operations. When compared to the microprocessor, the FPGA is over 100 times more power efficient. Similarly, the compute throughput per unit chip cost of FPGAs is 20% to 307% more than the 1 GHz DSP, and 50% to 505% more than the 3.8GHz Pentium-4 processor.

We are currently developing more advanced Radio Astronomy applications. By the end of summer 2005 (in time for the symposium), we expect to have an 8-antenna correlator system prototype running on the Green Bank Telescope dipole antenna array. We plan to develop a similar correlator for the Allen Telescope Array (ATA) with 32 antennas in the second half of 2005. For the final 350 antenna version of the ATA, 121 BEE2 modules will be employed providing an aggregate computational throughput of over 200 TeraOp/s.

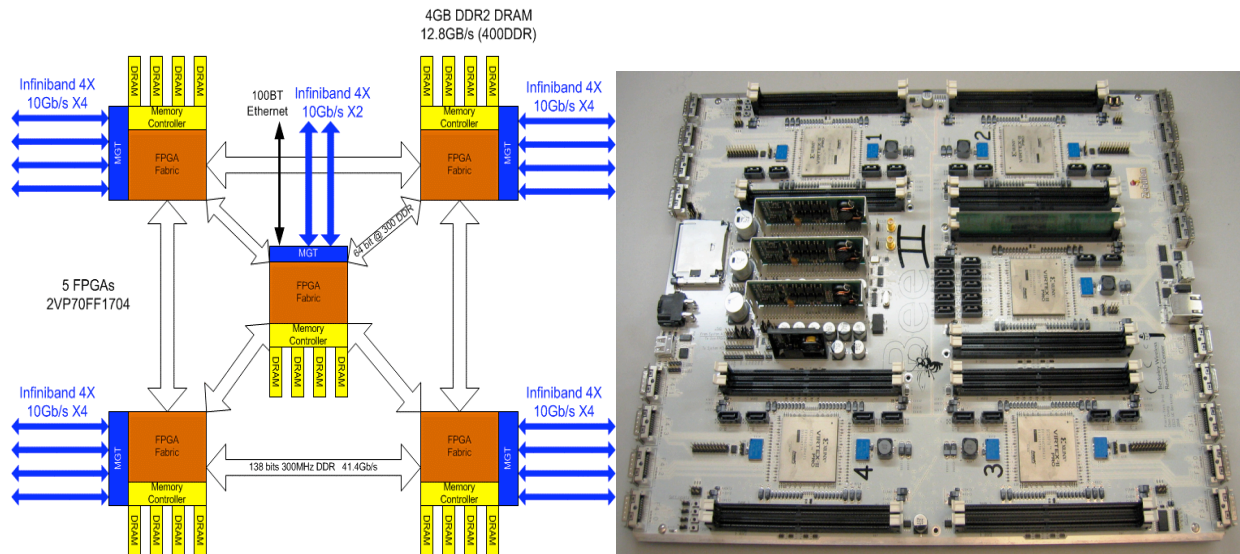


Figure 1: Compute Module Architecture Diagram and Photo